

How Information on Acoustic Scenes and Sound Events Mutually Benefits Event Detection and Scene Classification Tasks

Ami Igarashi* Keisuke Imoto* Yuka Komatsu* Shunsuke Tsubaki* Shuto Hario* Tatsuya Komatsu†

* Doshisha University, Japan

† LINE Corporation

Abstract—Acoustic scene classification (ASC) and sound event detection (SED) are fundamental tasks in environmental sound analysis, and many methods based on deep learning have been proposed. Considering that information on acoustic scenes and sound events helps SED and ASC mutually, some researchers have proposed a joint analysis of acoustic scenes and sound events by multitask learning (MTL). However, conventional works have not investigated in detail how acoustic scenes and sound events mutually benefit SED and ASC. We, therefore, investigate the impact of information on acoustic scenes and sound events on the performance of SED and ASC by using domain adversarial training based on a gradient reversal layer (GRL) or model training with fake labels. Experimental results obtained using the TUT Acoustic Scenes 2016/2017 and TUT Sound Events 2016/2017 show that pieces of information on acoustic scenes and sound events are effectively used to detect sound events and classify acoustic scenes, respectively. Moreover, upon comparing GRL- and fake-label-based methods with single-task-based ASC and SED methods, single-task-based methods are found to achieve better performance. This result implies that even when using single-task-based ASC and SED methods, information on acoustic scenes may be implicitly utilized for SED and vice versa.

I. INTRODUCTION

Environmental sound analysis [1], [2] is the analysis of audio recordings that are not limited to voice or music and has various real-world applications such as machine condition monitoring, automatic surveillance, media tagging, and biomonitoring systems [3], [4], [5], [6]. In environmental sound analysis, acoustic scene classification (ASC) and sound event detection (SED) are the fundamental research topics. In ASC, an acoustic scene label is estimated from a sound recording where the acoustic scene is defined as the place or situation in which the audio is recorded, such as *home*, *train*, or *meeting*. In SED, sound event labels and their timestamps in the sound recording are predicted, where a sound event is defined as a sound class, such as *bird singing*, *cutlery*, or *car horn*.

Recently, many systems for environmental sound analysis have been implemented using neural networks. For example, Valenti et al. [7], Liping et al. [8], Tanabe et al. [9], and Raveh and Amar [10] introduced ASC methods based on the convolutional neural network (CNN), Xception, VGG, and ResNet, respectively. Çakır et al. introduced the convolutional recurrent neural network (CRNN), which can capture temporal

and spectral information of sound events [11], for the SED task. Kong et al. [12] and Miyazaki et al. [13] proposed SED methods using a Transformer and Conformer encoder, respectively.

The conventional methods for environmental sound analysis address scene classification and event detection tasks separately. However, acoustic scenes and sound events are closely related; for example, in the acoustic scene *office*, the sound events *keyboard typing* and *mouse clicking* tend to occur, whereas the sound events *cutlery* and *car horn* occur infrequently, as shown in Fig. 1. Thus, the information on the sound events *keyboard typing* and *mouse clicking* will help in estimating the acoustic scene *office* and vice versa. Considering the relationship between acoustic scenes and sound events, Mesaros et al. [14] and Heittola et al. [15] proposed SED methods utilizing information on acoustic scenes. Imoto and co-workers proposed ASC methods based on Bayesian generative models, in which information on sound events is considered [16], [17]. Bear et al. [18], Tonami et al. [19], and Jung et al. [20] presented methods of jointly analyzing acoustic scenes and sound events based on the multitask learning (MTL) of ASC and SED. These works have revealed that utilizing the relationship between acoustic scenes and sound events improves the performance of each downstream task. However, conventional studies have not fully clarified how acoustic scenes and sound events mutually benefit event detection and scene classification tasks. In this work, we thus investigate the impact of information on acoustic scenes and sound events on the performance of SED and ASC by using domain adversarial training with a gradient reversal layer (GRL) [21] or model training with fake labels.

The remainder of this paper is structured as follows. In section 2, we overview the conventional scene classification and event detection methods based on single and multitask learning. Section 3 introduces the model training based on GRL and fake labels. In section 4, we discuss in detail how information on acoustic scenes and events affects the performance of ASC and SED, referring to the results of evaluation experiments. We conclude this paper in section 5.

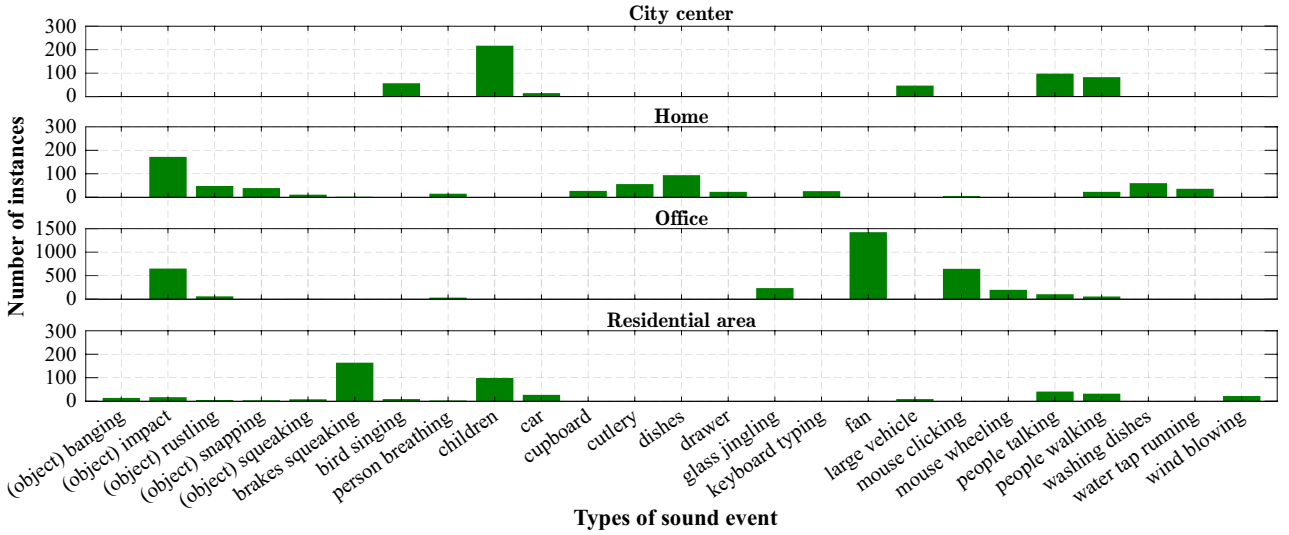


Fig. 1: Frequency of sound event instances for each acoustic scene in dataset used for evaluation experiments

II. CONVENTIONAL METHODS

A. Conventional Scene Classification and Event Detection Methods

In this section, we overview the conventional works of ASC and SED. Recently, many neural-network-based approaches such as CNN [7], [22], CRNN [11], and Transformer encoder [12] have been proposed. These methods use the time-frequency representation of the acoustic signal $X \in \mathcal{R}^{D \times T}$, such as the log mel-band spectrogram, as the acoustic feature, where D and T indicate the numbers of frequency bins and time frames, respectively. This acoustic feature is then fed to the ASC or SED network.

The model parameters of ASC are trained using the network output y_n and the cross-entropy (CE) loss function $\mathcal{L}_{\text{scene}}$,

$$\mathcal{L}_{\text{scene}} = -\sum_{n=1}^N \left\{ s_n \log(y_n) \right\}, \quad (1)$$

where N and s_n are the number of acoustic scene classes and the target scene label, respectively.

The model parameters of SED are trained using the network output $y_{t,m}$ and the following binary cross-entropy (BCE) loss function $\mathcal{L}_{\text{event}}$,

$$\begin{aligned} \mathcal{L}_{\text{event}} &= -\sum_{t=1}^T \left\{ \mathbf{z}_t \log(\mathbf{y}_t) + (1 - \mathbf{z}_t) \log(1 - \mathbf{y}_t) \right\} \\ &= -\sum_{t,m=1}^{T,M} \left\{ z_{t,m} \log(y_{t,m}) + (1 - z_{t,m}) \log(1 - y_{t,m}) \right\}, \end{aligned} \quad (2)$$

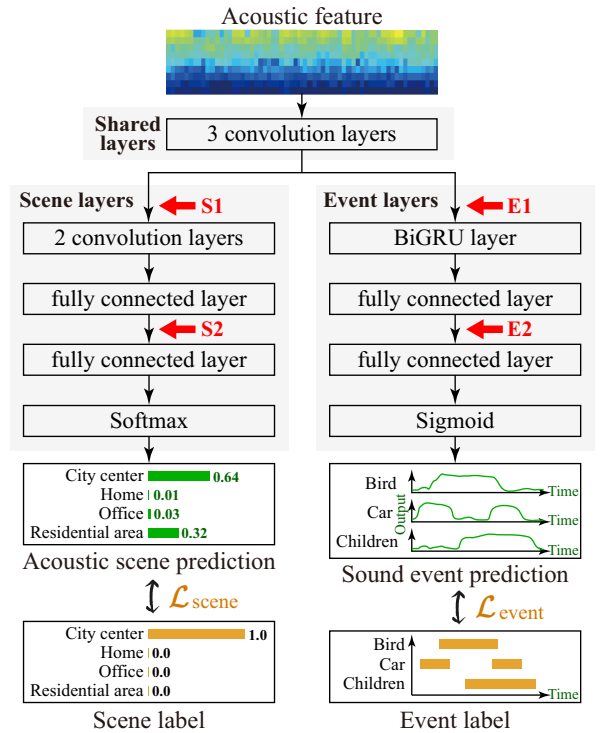


Fig. 2: Network structure of MTL-based joint analysis of ASC and SED [19]

where T , M , and $z_{t,m}$ indicate the number of time frames, the number of sound event classes, and the target event label, respectively.

B. Joint Analysis of Acoustic Scenes and Sound Events Based on Multitask Learning

Most conventional works address ASC and SED separately; however, many acoustic scenes and sound events are related mutually. Considering that knowledge of acoustic scenes and sound events can mutually aid in their estimation, the joint analysis of acoustic scenes and sound events based on multi-task learning has been proposed [18], [19], [23].

The network structure of MTL for the joint analysis of acoustic scenes and sound events is shown in Fig. 2. The MTL network consists of shared, scene-specific, and event-specific layers. The shared layers have convolution layers for extracting the embedded acoustic features and are expected to extract the information common to acoustic scenes and sound events. The scene- and event-specific layers include the CNN and recurrent neural network (RNN) layers, which are used for the downstream tasks of ASC and SED.

To train the MTL model of ASC and SED, the following loss function \mathcal{L} is applied:

$$\mathcal{L} = \alpha \mathcal{L}_{\text{scene}} + \beta \mathcal{L}_{\text{event}}, \tag{3}$$

where α and β are the constant weighting factors of ASC and SED losses, respectively. In this work, $\beta = 1.0$ can be set without loss of generality.

III. METHODOLOGY

Some conventional studies on jointly analyzing acoustic scenes and sound events revealed that pieces of information on acoustic scenes and sound events improve the performance of SED and ASC, respectively [19]. However, how acoustic scenes and sound events benefit SED and ASC tasks mutually has not been fully investigated in conventional studies. To evaluate how information on acoustic scenes and sound events benefits the performance of ASC and SED, we apply the domain adversarial training based on GLR and fake labels of acoustic scenes and sound events, which enables the training of the model without intentionally using information on acoustic scenes and sound events.

A. MTL Based on Domain Adversarial Training

We first apply the domain adversarial training based on GLR [21] to the MTL of ASC and SED. The GRL acts as an identity transformation of the input during the forward propagation, but changes the sign of input, i.e., it multiplies by -1 during the backpropagation as follows.

$$\text{Forward: } G(\mathbf{x}) = \mathbf{x} \tag{4}$$

$$\text{Backward: } \frac{dG}{dx} = -\lambda \mathbf{I} \tag{5}$$

Here \mathbf{x} , λ , and \mathbf{I} are the input of the GRL, the weighting factor, and the identity matrix, respectively. By adding the GRL to a network, the model parameters in layers prior to

the GRL are trained as increasing the prediction error. Thus, the GRL enables the training of models that do not depend on the downstream task with this layer. In the experiments conducted in this work, we added the GRL to either position shown by the red arrows (S1, S2, E1, and E2) in Fig. 2 and evaluated the performance of ASC and SED.

B. MTL Based on Fake Label

We also evaluate the performance of ASC and SED using the MTL model trained with fake labels. In this work, we created fake scene and event labels as follows.

$$\text{Fake scene label: } \hat{\mathbf{s}} = \text{shuffle}(\mathbf{s}) \tag{6}$$

$$\text{Fake event label: } \hat{\mathbf{z}}_t = \text{shuffle}(\mathbf{z}_t) \tag{7}$$

Here, shuffle, \mathbf{s} , and \mathbf{z}_t indicate the random reordering operation of vector elements, the target scene label vector, and the target event label vector in time frame t , respectively. In the experiment, we replaced either the acoustic scene or sound event labels with the fake labels. As with the GRL, the model training with fake labels allows us to investigate how acoustic scenes and sound events affect ASC and SED.

IV. EVALUATION EXPERIMENTS

A. Experimental Conditions

We conducted evaluation experiments to investigate how the information on acoustic scenes and sound events can benefit ASC and SED tasks. In the evaluation experiments, we used a dataset composed of parts of the TUT Acoustic Scenes 2016/2017 and TUT Sound Events 2016/2017 [24], [25]. From these four datasets, we selected sound clips of four acoustic scenes, *city center*, *home*, *office*, and *residential area*, which include 25 sound event classes. The data comprises a total of 266 min of sounds (192 min for the development set and 74 min for the evaluation set). The details of the dataset used for the evaluation experiments are found in [26].

As an acoustic feature, we extracted the 64-dimensional log mel-band energy, which was calculated with a 40 ms frame length and a 20 ms hop size. We used the same MTL network structure as that of the conventional method [19], as shown in Table I, where BiGRU and FC layers indicate the bidirectional gated recurrent unit and fully connected layers, respectively. For each method, we conducted the evaluation experiments 10 times with random initial values of model parameters. Other experimental conditions are listed in Table II.

B. Experimental Results

1) *Overall Performance of ASC and SED:* Table III shows the overall performance of ASC and SED. CNN (ASC) and CNN-BiGRU (SED) are the single task networks for ASC and SED; they have the same network structures as the shared + scene layers and shared + event layers in Table I, respectively. GRL (S1) indicates the MTL-based method with GRL in S1 of Fig. 2. In this experiment, we evaluated the SED system

TABLE I: Details of structure of MTL network of ASC and SED

Shared layers	
Log-mel energy	
500 frames × 64 mel bin	
3×3 kernel size / 128 ch.	
Batch norm., Leaky ReLU	
1×8 Max pooling	
$\left(\begin{matrix} 3 \times 3 \text{ kernel size / 128 ch.} \\ \text{Batch norm., Leaky ReLU} \\ 1 \times 2 \text{ Max pooling} \end{matrix} \right) \times 2$	
Scene layers	Event layers
3×3 kernel size / 256 ch.	
Batch norm., Leaky ReLU	BiGRU w/ 32 units
25×1 Max pooling	
3×3 kernel size / 256 ch.	
Batch norm., Leaky ReLU	FC w/ 32 units, Leaky ReLU
Global max pooling	
FC w/ 32 units, Leaky ReLU	FC w/ 25 units, Sigmoid
FC w/ 4 units, Softmax	

TABLE II: Experimental conditions

Acoustic feature	Log-mel energy (64 dim.)
Frame length / shift	40 ms / 20 ms
Length of sound clip	10 s
Optimizer	Adam [27]
α, β	0.0001, 1.0
λ	1.0

with the frame-based metric because the SED system outputs the frame-wise predictions, and we can understand the basic trends of the system output with this metric.

The results show that when we apply the GRLs in the scene and event layers or train models with fake labels, the performance of ASC and SED tends to decrease. These results indicate that pieces of information on acoustic scenes and sound events are effectively used to detect sound events and classify acoustic scenes, respectively. A comparison of GRL- and fake-label-based methods with single-task-based ASC and SED methods reveals that single-task-based methods achieve better performance. This result implies that even when using single-task-based ASC and SED methods, information on acoustic scenes may be implicitly utilized for SED and vice versa.

2) *Details of Performance of SED and ASC:* To investigate the event detection performance in detail, we show the Fscores for selected sound events in Table IV. Note that a similar trend is observed for sound events not shown in Table IV. This result shows that the Fscores decrease when we train the models with GRL (S1), GRL (S2), and fake scene labels. This result also supports that information on acoustic scenes is effectively used to detect sound events. Moreover, the experimental result shows that CNN-BiGRU, which is the single-task-based method, outperforms GRL (S1), GRL (S2), and fake scene labels. Thus, when detecting most sound events using the single-task-based method, information on acoustic

TABLE III: Overall performance of ASC and SED

Method	Scene		Event	
	Micro-Fscore	Macro-Fscore	Micro-Fscore	Macro-Fscore
CNN (ASC)	85.00%	84.29%	-	-
CNN-BiGRU (SED)	-	-	42.54%	11.09%
Conventional MTL	88.57%	88.85%	44.63%	11.57%
GRL (S1)	77.89%	76.14%	40.08%	9.64%
GRL (S2)	33.34%	22.03%	39.59%	9.92%
Fake scene labels	24.45%	11.55%	41.57%	10.43%
GRL (E1)	60.90%	56.30%	13.07%	1.12%
GRL (E2)	75.23%	73.99%	0.02%	0.02%
Fake event labels	84.22%	84.09%	0.00%	0.00%

TABLE IV: Average Fscores for selected sound events

Method	bird singing	car	dishes	keyboard typing	large vehicle
CNN-BiGRU (SED)	46.36%	44.53%	0.17%	4.34%	12.27%
Conv. MTL	46.29%	45.51%	0.25%	5.08%	12.29%
GRL (S1)	40.96%	43.46%	0.00%	0.56%	9.28%
GRL (S2)	37.62%	40.12%	0.01%	1.21%	10.54%
Fake scene label	28.93%	39.96%	0.06%	1.63%	12.05%

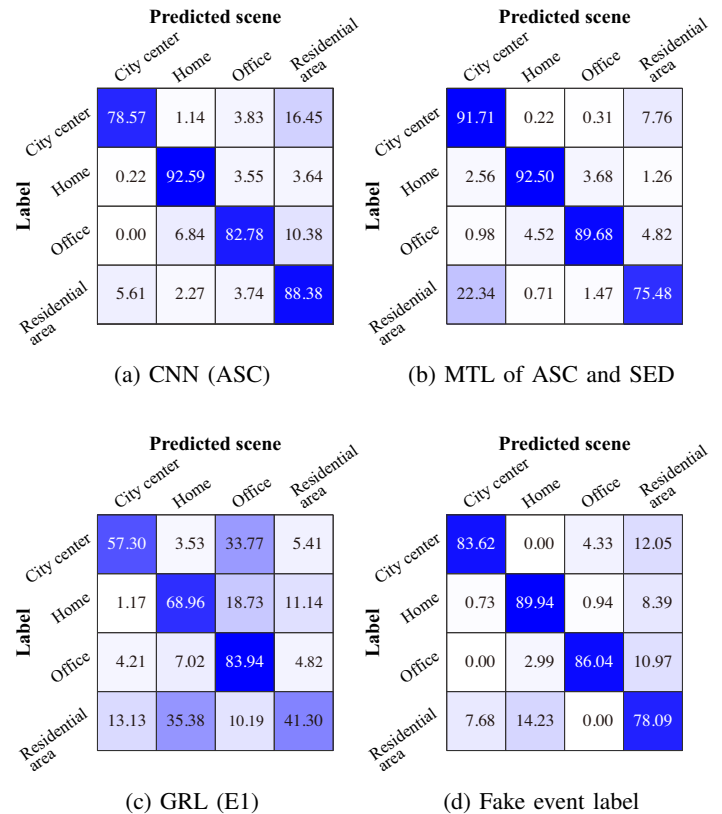


Fig. 3: Confusion matrix for each method in terms of recall (%)

scenes may also benefit SED implicitly.

The detailed classification results of acoustic scenes are shown in Fig. 3. Compared with CNN (ASC), the MTL of ASC and SED has fewer misclassifications of *office*. Moreover, the

MTL of ASC and SED still confuses *city center* and *residential area*. This is because in *city center* and *residential area*, similar sound events tend to occur, as shown in Fig. 1; thus, providing information on sound events yields fewer benefits in distinguishing between these acoustic scenes. The comparison of CNN (ASC) and GRL (E2) also shows that information on annotated sound events is used in the single-task-based ASC method implicitly. On the other hand, the impact of unlabeled sound events, such as background noise, on ASC must be further investigated in future work.

V. CONCLUSIONS

In this study, we investigated how information on acoustic scenes and sound events benefits SED and ASC in detail. To evaluate this, we applied domain adversarial training based on GRL and fake labels of acoustic scenes and sound events. The evaluation experiments conducted using parts of the TUT Acoustic Scenes 2016/2017 and TUT Sound Events 2016/2017 datasets indicated that pieces of information on acoustic scenes and sound events were effectively utilized for SED and ASC, respectively. Moreover, by comparing GRL- and fake-label-based methods with single-task-based methods, we found that single-task-based methods achieve better performance. This implies that even when using single-task-based ASC and SED methods, information on acoustic scenes may be implicitly utilized for SED and vice versa.

ACKNOWLEDGMENT

This work was supported by JSPS KAKENHI Grant Number JP20H00613.

REFERENCES

- [1] T. Virtanen, M. Plumbley, and D. Ellis, Eds., *Computational Analysis of Sound Scenes and Events*. Springer, 2017.
- [2] K. Imoto, "Introduction to acoustic event and scene analysis," *Acoustical Science and Technology*, vol. 39, no. 3, pp. 182–188, 2018.
- [3] Y. Koizumi, Y. Kawaguchi, K. Imoto, T. Nakamura, Y. Nikaido, R. Tanabe, H. Purohit, K. Suefusa, T. Endo, M. Yasuda, and N. Harada, "Description and discussion on DCASE2020 challenge task2: Unsupervised anomalous sound detection for machine condition monitoring," *Proc. Detection and Classification of Acoustic Scenes and Events (DCASE)*, pp. 81–85, 2020.
- [4] C. Chan and E. W. M. Yu, "An abnormal sound detection and classification system for surveillance applications," *Proc. European Signal Processing Conference (EUSIPCO)*, pp. 1851–1855, 2010.
- [5] E. Fonseca, M. Plakal, F. Font, D. P. W. Ellis, X. Favory, J. Jordi, and X. Serra, "General-purpose tagging of freesound audio with AudioSet labels: Task description, dataset, and baseline," *Proc. Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, pp. 69–73, 2018.
- [6] V. Morfi, R. F. Lachlan, and D. Stowell, "Deep perceptual embeddings for unlabelled animal sound," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 150, no. 1, pp. 2–11, 2020.
- [7] M. Valenti, S. Squartini, A. Diment, G. Parascandolo, and T. Virtanen, "A convolutional neural network approach for acoustic scene classification," *Proc. International Joint Conference on Neural Networks (IJCNN)*, pp. 1547–1554, 2017.
- [8] Y. Liping, C. Xinxing, and T. Lianjie, "Acoustic scene classification using multi-scale features," *Proc. Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, pp. 29–33, 2018.
- [9] R. Tanabe, T. Endo, Y. Nikaido, T. Ichige, P. Nguyen, Y. Kawaguchi, and K. Hamada, "Multichannel acoustic scene classification by blind dereverberation, blind source separation, data augmentation, and model ensembling," *Tech. Rep. DCASE Challenge 2018 Task5*, pp. 1–4, 2018.
- [10] A. Raveh and A. Amar, "Multi-channel audio classification with neural network using scattering transform," *Tech. Rep. DCASE Challenge 2018 Task5*, pp. 1–4, 2018.
- [11] E. Çakır, G. Parascandolo, T. Heittola, H. Huttunen, and T. Virtanen, "Convolutional recurrent neural networks for polyphonic sound event detection," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 25, no. 6, pp. 1291–1303, 2017.
- [12] Q. Kong, Y. Xu, W. Wang, and M. D. Plumbley, "Sound event detection of weakly labelled data with CNN-Transformer and automatic threshold optimization," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 28, pp. 2450–2460, 2020.
- [13] K. Miyazaki, T. Komatsu, T. Hayashi, S. Watanabe, T. Toda, and K. Takeda, "Convolution-augmented transformer for semi-supervised sound event detection," *Tech. Rep. DCASE Challenge 2020 Task4*, pp. 1–4, 2020.
- [14] A. Mesaros, T. Heittola, and A. Klapuri, "Latent semantic analysis in sound event detection," *Proc. European Signal Processing Conference (EUSIPCO)*, pp. 1307–1311, 2011.
- [15] T. Heittola, A. Mesaros, A. Eronen, and T. Virtanen, "Context-dependent sound event detection," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2013, no. 1, 2013.
- [16] K. Imoto and S. Shimauchi, "Acoustic scene analysis based on hierarchical generative model of acoustic event sequence," *IEICE Transactions on Information and Systems*, vol. E99-D, no. 10, pp. 2539–2549, 2016.
- [17] K. Imoto and N. Ono, "Acoustic topic model for scene analysis with intermittently missing observations," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 27, no. 2, pp. 367–382, 2019.
- [18] H. L. Bear, I. Nolasco, and E. Benetos, "Towards joint sound scene and polyphonic sound event recognition," *INTERSPEECH*, pp. 4594–4598, 2019.
- [19] N. Tonami, K. Imoto, R. Yamanishi, and Y. Yamashita, "Joint analysis of sound events and acoustic scenes using multitask learning," *IEICE Transactions on Information and Systems*, vol. E104-D, no. 02, pp. 294–301, 2021.
- [20] J. Jung, H. J. Shim, J. H. Kim, and H. J. Yu, "DCASENET: An integrated pretrained deep neural network for detecting and classifying acoustic scenes and events," *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 621–625, 2021.
- [21] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. March, and V. Lempitsky, "Domain-adversarial training of neural networks," *Journal of Machine Learning Research*, vol. 17, no. 59, pp. 1–35, 2016.
- [22] S. Hershey, S. Chaudhuri, D. P. W. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold, M. Slaney, R. J. Weiss, and K. Wilson, "CNN architectures for large-scale audio classification," *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 131–135, 2017.
- [23] K. Imoto, N. Tonami, Y. Koizumi, M. Yasuda, R. Yamanishi, and Y. Yamashita, "Sound event detection by multitask learning of sound events and scenes with soft scene labels," *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 621–625, 2020.
- [24] A. Mesaros, T. Heittola, and T. Virtanen, "TUT database for acoustic scene classification and sound event detection," *Proc. European Signal Processing Conference (EUSIPCO)*, pp. 1128–1132, 2016.
- [25] A. Mesaros, T. Heittola, A. Diment, B. Elizalde, A. Shah, B. Raj, and T. Virtanen, "DCASE 2017 challenge setup: Tasks, datasets and baseline system," *Proc. Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, pp. 85–92, 2017.
- [26] <https://www.ksuke.net/dataset>.
- [27] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han, "On the variance of the adaptive learning rate and beyond," *Proc. International Conference on Learning Representations (ICLR)*, pp. 1–13, 2020.